

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**A Dynamically Generated Schema Representing
Multiple Hierarchies of Inter-Object Relationships**

Inventor(s):

Kim Cameron

Stewart MacLeod

George G. Robertson

James Booth

Luc Clément

ATTORNEY'S DOCKET NO. MS1-776US

TECHNICAL FIELD

The following subject matter relates to object-to-object or "inter-object" relationships. More particularly, the subject matter pertains to interfacing with dynamically generated multiple hierarchies of inter-object relationships based on the values of attributes of the objects.

BACKGROUND

Any object can be linked, correlated, associated, differentiated, or in some manner categorized with respect to a different object to form implicit or explicit inter-object relationships. For instance, in an organization, a person typically has implicit and explicit relationships with other people in the organization, organizational resources (e.g., printers, facilities, etc.), geographical locations, business units, club memberships, and so on. Each implicit and/or explicit relationship between respective objects (i.e., the person, the other people, a resource, etc.) represents a respective hierarchical data relationship.

For example, one hierarchical data relationship is represented by each person within the company that has access to a specific resource (e.g., a building on the company campus, a room, a printer, etc); the resource being the root node of the hierarchy and the individuals with access to the resource being the leaves. Another hierarchical data relationship is represented by individuals that make up the management structure of the company. Other inter-object data relationships may represent a hierarchy of individuals within a particular business unit, all employees of the company that have specialized training, and so on.

1 Unfortunately, even though a data store can be configured to some extent
2 by a network administrator to represent inter-object relationships within
3 hierarchies of other data, complex inter-object relationships (e.g., such as those
4 representing a single object within more ^{THAN} that one hierarchy) are not simply and
5 adequately represented using conventional data store (e.g., directory, database,
6 etc.) systems and technologies. (Traditional directories include those based on the
7 well-known X.500 standard and the Lightweight Directory Access Protocol
8 (LDAP).

9 To illustrate this limitation of traditional data store systems and
10 technologies, consider that a directory typically represents inter-object
11 relationships using rigid data naming and inflexible directory schemas. Objects or
12 nodes in the directory are organized within a single hierarchy with a root node at
13 the top of the hierarchy. The root node has a name. Each other node in the
14 directory is named based on its direct naming relationship to the root node and
15 also with respect to each intervening node in the respective node's hierarchy. As a
16 result, if a parent object is renamed in a single operation, any objects that are
17 subordinate or children of the parent object are also renamed in that same single
18 operation. This is because an object's full "distinguished name" includes the name
19 of each parent object(s) all the way down the line to the root node's name.

20 It is the full distinguished name of an object that also represents its static
21 location or data relationship with respect to each other object in the data store.
22 Thus, an object's distinguished name inflexibly inter-tangles object naming within
23 a single hierarchy with inter-object relationships in that hierarchy. Because of this,
24 any navigation of the data store must be performed from top-to-bottom to
25

1 determine and subsequently present any inter-object relationships—that is from
2 the root object, to a parent object to any subordinate child object(s).

3 Because traditional data stores (e.g., directories, databases, and so on) rely
4 on a carefully specified and inflexible object naming scheme to identify inter-
5 object relationships, an administrator configuring the data store requires a-priori
6 knowledge of the inter-object relationships when configuring the data store.
7 Additionally, any configuration of the data store must consider not only the proper
8 representation of inter-object relationships in the data store, but must also consider
9 the heuristics that a search engine requires to navigate the data store.

10 To make matters worse, elastic data relationships are not easily described,
11 represented, or navigated using conventional data store systems and techniques.
12 An elastic data relationship is one wherein the relationship is derived from data
13 that defines an object at any point in time. This means that over time elastic data
14 relationships can be dynamic. For instance, consider the following non-obvious
15 and potentially elastic data relationships: a Web site and the Web pages that make
16 up the Web site, a customer and the individuated services that the customer
17 purchases from a merchant, a personal computer (PC) and peripheral devices that
18 are coupled to the PC, a city and the districts within the city, a business and the
19 business' contacts, an employee and the employee's dependents, and the like.

20 These non-obvious and potentially elastic data relationships are not easily
21 represented because whenever a one-to-one correspondence between a surface
22 object and corresponding sub-objects needs to be represented in the data store, an
23 irreversible design choice must be made. (Conventional practice is to strictly
24 control directory schema updates due to the serious nature of directory schema
25 modification). A network administrator can opt for “total incorporation” of the

sub-objects into the particular object by representing the sub-objects as attributes of the surface object in the directory schema. Or the network administrator can opt for "total distinction" of each object, by creating separate objects in the schema for sub-object components, and positioning the separate objects subordinate to the surface object.

To illustrate this irreversible design choice, consider that a particular network router includes multiple router modules plugged into the router's backplane. Information about the router and the router modules are typically stored in a directory in one or two different fashions—each of which may be equally unsatisfactory depending on how entities and their respective relationships to other entities are represented. One design choice is to characterize a router and its corresponding router modules as a single hierarchical data structure representing the network router as a parent object, and the corresponding router modules as child objects that are subordinate to the parent object. A different design choice is to characterize the router and the router's associated router modules as a single parent object with complex attributes. The parent object represents the router (backplane), and the complex attributes representing the respective router modules that are hosted by the router.

In consideration of the first design choice, depending how the router and the modules are configured, collapsing information about the router modules, or boards onto the backplane may prove unwieldy. This is because the functionality of the router's backplane may be small as compared to the functionality of the network router modules hosted by the router. Whereas considering the second design choice, completely separating the boards from the backplane may be

1 equally unsatisfactory because the router is still a single physical router box that
2 generally includes a number of router modules. ↙

3 Both of the described solutions to representing data relationships with an
4 inflexible directory schema are time consuming to implement and counter-
5 intuitive. The semantics of shape and naming in the directory must be agreed on
6 in advance to solve the simplest design problem. Thus, whenever a one-to-one
7 correspondence between an entity and corresponding sub-entities needs to be
8 represented in a traditional directory, an irreversible and inflexible design choice
9 must be made within the directory schema.

10 Whichever design choice is selected, the data store and tools used to
11 navigate, search and present objects within the data store with respect to inter-
12 object relationships have been substantially limited. This is because the data store
13 itself can not represent all of the possible implicit and explicit inter-object
14 relationships of an object. This is considered by many computer programmers to
15 be one of the most intractable problems of directory schema in traditional
16 directories. This is also deemed to be the reason that computer program
17 applications are not typically portable across directory platforms or even directory
18 instances.

19 To further worsen matters, recent developments in information technology
20 provide network administrators with opportunities to tie disparate data stores (e.g.,
21 databases, directories, and so on) of data together into a single logical directory or
22 "metadirectory". Such disparate databases and directories include, for example,
23 information corresponding to enterprise users, tangible and intangible resources,
24 financial information, corporate e-mail systems, network operating systems, and
25 so on.

Metadirectories present network administrators with complex and often elastic object data relationships that cannot be simply or adequately described, represented, navigated, or presented using traditional systems and procedures to configure and manage data stores. Considerable efforts are required on the part of the administrator (or a staff of administrators) to configure a data store. Manually determining and implementing such inter-object relationships (whether they be dynamic or not) is fraught with the potential for human error and oversight. Furthermore, database administrators with an appropriate level of such knowledge to perform such a directory configuration are expensive.

The following described subject matter addresses these and other problems of representing inter-object relationships.

SUMMARY

The described arrangements and procedures provide for interfacing (e.g., managing, presenting, etc.) with complex and often elastic inter-object relationships between objects in a data polyarchy. Specifically, a schema is dynamically generated by a server to represent multiple hierarchies of inter-object relationships between objects in a data polyarchy. The schema indicates or lists each attribute or element of interest in the data polyarchy. The schema further indicates any of one or more dimensions of inter-object relationships within which objects that comprise at least a subset of the listed the elements of interest participate. Thus, the schema indicates how to interface with the data polyarchy, which represents multiple hierarchies of inter-object relationships based on the values of attributes of the represented objects.

1 The generated schema is communicated to a client by a data polyarchy
2 server.

3 Responsive to receiving a request from the client based on the schema, the
4 server accesses one or more objects in the data polyarchy based on the request.
5 The server transforms the one or more of the objects into data that expresses any
6 inter-object relationships between the one or more objects in the data polyarchy
7 based on the request. The server issues the transformed data to the requesting
8 client.

9 **BRIEF DESCRIPTION OF THE DRAWINGS**

10
11 Fig. 1 shows an exemplary system for dynamically generating and
12 managing multiple hierarchies of inter-object relationships based on the values of
13 attributes of the objects.

14 Fig. 2 illustrates an exemplary polyarchy data structure to represent
15 multiple hierarchies of dynamically generated inter-object relationships that are
16 based on the values of attributes of the objects.

17 Fig. 3 shows an exemplary schema data structure to indicate how a data
18 polyarchy of Fig. 2 can be created, accessed, and manipulated in a meaningful
19 manner.

20 Fig. 4 shows an exemplary procedure to generate multiple hierarchies of
21 inter-object relationships based on the values of attributes of the objects.

22 Fig. 5 shows an exemplary polyarchical query language (PQL) request used
23 by a client to request a server to return information (a PQL response) from a data
24 polyarchy.
25

Fig. 6 shows a user interface (UI) displaying an exemplary PQL query and a corresponding exemplary PQL response. Specifically, the PQL query includes a modifier parameter based on a data polyarchy schema to specify a particular attribute with which to perform a search of a polyarchical data set.

Fig. 7 is a block diagram of a UI displaying an exemplary PQL query and a corresponding exemplary PQL response. Specifically, the PQL query includes a modifier parameter based on an elements-of-interest schema; the parameter specifies a limiting attribute with which to modify a result of a search.

Fig. 8 is a block diagram of a UI displaying an exemplary PQL query and a corresponding exemplary PQL response. Specifically, the PQL query includes a Boolean modifier parameter to perform a mathematical operation with respect to polyarchies of data relationships.

Fig. 9 is a block diagram of a UI showing an exemplary PQL query and a corresponding exemplary PQL response. Specifically, the PQL query includes a dimension information indicator parameter for specifying a dimension within which to view an object stored in a data store.

Fig. 10 is a UI showing an exemplary PQL query and a corresponding exemplary PQL response. Specifically, the PQL query includes a dimension information modifier parameter, which specifies a particular hierarchical direction and a particular hierarchical depth for a server process to present a data relationship between a complex object of a polyarchical data set and one or more other objects.

Fig. 11 is a UI showing use of a locating element in an exemplary PQL query with respect to a particular attribute and a subsequent intersection between

two corresponding polyarchies of data relationships to form an exemplary PQL response.

Fig. 12 is UI showing an exemplary PQL query and a corresponding exemplary PQL response. Specifically, the PQL query illustrates use of filter and union parameters with respect to two polyarchies of data relationships.

Fig. 13 illustrates aspects of an exemplary procedure to manage data (e.g., to access, present, provide, and/or manipulate objects, etc.) in a data polyarchy (i.e., multiple hierarchies of dynamically generated inter-object relationships that are based on the values of attributes of the objects).

Fig. 14 shows aspects of an exemplary operating environment for managing a data polyarchy.

DETAILED DESCRIPTION

Overview

The following subject matter replaces traditional notions of complex real-world object presentation within a single static hierarchy, wherein directory object naming and inter-object relationships are inter-tangled and unwieldy for representing complex data relationships. More specifically, traditional notions of distinguished names for representing inter-object relationships within a single directory of static inter-object relationships are replaced with graphs of elastic (non-static) inter-object connections in multiple dimensions of data relationships (e.g., mono and/or bi-directional relationships) based on attributes of the objects. In other words, the data relationships establish that one or more data objects participate in one or more respective dimensions, or polyarchies of inter-object

relationships. One or more of these hierarchies can intersect creating intersecting hierarchies of inter-object relationships.

Dynamically generated multiple hierarchies of data relationships based on object attributes are represented in a data polyarchy. Specifically, the data polyarchy is generated using each object's respective data attributes or data values and multifarious interrelationships of those values with attributes that correspond to other objects in the polyarchical data set. The inter-object relationships in the data polyarchy can be elastic because inter-object relationships are derived from data defined by an object at any point in time. Patterns of relationships between objects emerge by presenting an object in one or more "dimensions" or polyarchies of data relationships. Such relationships are presented using inter-object connections between virtual entities representing the objects. A virtual entity corresponds to an object of interest and includes and organizes information about an object of interest—including information about how to get more information about the object of interest. Such objects can be presented to people or computer programs that embody that interest.

In contrast to traditional systems and procedures to represent inter-object relationships in a data store, the following described arrangements and procedures are dynamic, in that they are automated and do not require any manual intervention from a network administrator to configure inter-object relationships. By dynamically generating a data polyarchy complex inter-object relationships based on object data are automatically determined without presenting any inflexible design choice to a schema designer.

This means that the network administrator or computer program (e.g., a search engine) is not required to have any a-priori knowledge of complex inter-

1 object relationships to generate, navigate, or search a data store. This also means
2 that each object in a data store can be viewed from as many different dimensional
3 inter-object hierarchies as apply to the respective object. Furthermore, as an
4 object's elastic data relationships change, the data polyarchy automatically detects
5 and reflects those changes.

6 The following description sets forth arrangements and procedures based on
7 a directory schema for representing polyarchies of inter-object relationships that
8 incorporates elements recited in the appended claims. The subject matter is
9 described with specificity to meet statutory requirements. However, the
10 description itself is not intended to limit the scope of this patent. Rather, the
11 inventors have contemplated that the claimed subject matter might also be
12 embodied in other ways, to include different elements or combinations of elements
13 similar to the ones described in this document, in conjunction with other present or
14 future technologies.

15 Exemplary System

16 Fig. 1 shows an exemplary system 100 to dynamically generate and manage
17 multiple hierarchies of inter-object relationships based on the values of attributes
18 of the objects. The system represents a distributed computing environment
19 including a data polyarchy server 102 operatively coupled across a network 104
20 to one or more other optional data servers 106, one or more databases 108, and
21 one or more client computers 110. The operative coupling of the data polyarchy
22 server to the network can be made in any number of different ways such as
23 through one or more server appliances (e.g., a server appliance on the outside of a
24
25

1 Web farm-server farm), a corporate portal (intranet), a local area network (LAN),
2 a co-located data store (e.g., a database 108), and so on.

3 The data polyarchy server 102 includes a processor 112 operatively coupled
4 to a memory 114 that includes computer-executable instructions 116 and data 118.
5 The processor is configured to fetch and execute the computer-executable
6 instructions and fetch the data during such execution. Such computer-executable
7 instructions include an operating system (not shown), and a data polyarchy
8 management module 120 to dynamically generate and manage multiple
9 hierarchies of inter-object relationships based on the values of attributes of the
10 objects. These dynamically generated multiple hierarchies of inter-object
11 relationships are stored in the polyarchical data set 122, which is also referred to
12 as the data polyarchy. To generate the data polyarchy 118, the data polyarchy
13 management module 120 uses data (e.g., Extensible Markup Language (XML)
14 data) from any number of different data sources such as from one or more other
15 optional servers 106 and/or databases 108. For instance, a server 106 provides
16 data (e.g., directories of enterprise users, resources, financial information,
17 corporate e-mail systems, network operating systems, etc.) to the data polyarchy
18 server from any number of various data stores—databases, directories,
19 metadirectories, and so on. A database 108 is a structured or unstructured data
20 store, including an object-oriented database such as an XML database, a Hypertext
21 Markup Language (HTML) database, an SQL server database, and so on.

22 Responsive to generating and managing the data polyarchy 122, the
23 management module 120 respectively generates and updates the elements of
24 interest schema 124. The elements-of-interest schema indicates how an optional
25

1 client computer 110 can manipulate and display the objects in the data polyarchy
2 with respect to their respective polyarchies of inter-object relationships.

3 For instance, the elements-of-interest schema 124 identifies each object in
4 the data polyarchy 122 as an address referencing a virtual entity (e.g., see the
5 virtual object 210 of Fig. 2) that represents the respective object. These virtual
6 entities are stored as vectors or arrays of addresses in the schema. Each different
7 type of attribute that an object in the data polyarchy could have is also identified in
8 the schema as well as what kinds of indexes are to be used on the various attribute
9 types. (A data index provides for object access). For each attribute type it is
10 convenient to store with its definition, its corresponding index. In this manner, for
11 example, if somebody requests for an attribute, the index is readily available and
12 all of the values assumed by the attribute can be determined very quickly. (An
13 elements-of-interest schema is described in greater detail below in reference to
14 Fig. 3).

15 The data polyarchy server 102 can generate any number of schemas 124.
16 Each generated schema can provide access to various subsets of the objects in the
17 data polyarchy 122 independent of the objects represented by other schemas 124.
18 For example, a first schema 124 can be distributed to network administrators to
19 provide access to resources and attributes such as printers and access lists that are
20 otherwise protected or hidden from other employees. In the same manner, a
21 second schema can be distributed to the president of human resources. While the
22 second schema may provide the president with access to certain privileged
23 employee records, the second schema could be completely silent with respect to
24 the resources that are available to the network administrators group via the first
25

1 schema. In this manner, schemas 124 can be designed to provide access control to
2 organizational resources.

3 The data polyarchy server 102 communicates the elements-of-interest
4 schema 124 to one or more optional clients 110. The client computer supports a
5 graphical user interface (not shown) for displaying inter-object relationships in the
6 data polyarchy 122 as described by the elements of interest schema. Exemplary
7 arrangements and procedures to display objects within polyarchies of data
8 relationships are described in related U.S. Patent application serial no.
9 09/728,935, titled "Hierarchy Polyarchy Visualization", filed on 11/29/00, which
10 is assigned to the assignee hereof, and which is incorporated by reference.

11 The data polyarchy 122 and the elements of interest schema 124 can be
12 replicated one or more times in a memory cache 114 by the data polyarchy server
13 102. An exemplary memory cache is described in greater detail below in
14 reference to Fig. 14. Since the polyarchical server can operate either data set from
15 a corresponding memory cache, there can be as many copies of the respective data
16 sets as necessary. Thus no matter how demanding a client 110, the data polyarchy
17 server can satisfy the demand.

18 When data polyarchy 122 and the elements of interest schema 124 are
19 replicated in a memory cache 114 by the data polyarchy server 102, the server can
20 maintain an authoritative store (not shown) in the memory 114 to represent the
21 most recent, or current representation of the inter-object relationships. Such an
22 authoritative store is beneficial because caches by their very nature are always out
23 of date to some degree—meaning that data in a cache is only as “fresh”, or timely
24 as the most recent cache update. In light of this, a client requesting information
25 from the data polyarchy 122 can indicate the level of data reliability or timeliness

1 required by the client. If a high timeliness is required, the server 102 can access
2 the data polyarchy from the authoritative store, rather than from more out of data
3 caches. The speed of access to an authoritative cache depends on its respective
4 implementation (e.g., implemented in internally to the server in random access
5 memory or externally to the server in a data storage device).

6 **Exemplary Data Polyarchy**

7 Fig. 2 shows an exemplary polyarchical data set 122 to represent multiple
8 dimensions of inter-object relationships based on attributes within the data. The
9 data is anything that can be differentiated (e.g., anything that is an object of
10 interest represented in a directory, database, etc., can be an object). The data set
11 122 is formatted to allow designers to create their own customized tags, enabling
12 the definition, transmission, validation, and interpretation of data between
13 applications and between organizations. For example, the data format can be an
14 XML data format.

15 The data polyarchy 122 includes multiple virtual object data fields 210.
16 Each virtual object data field includes and organizes information about a
17 respective object, including, for example, information about how to get more
18 information about the respective object. Specifically, the virtual object includes a
19 globally unique identifier (GUID) data field 212 and if appropriate for the
20 particular object, one or more attribute data fields 214.

21 A GUID 212 uniquely identifies the virtual object (which in turn represents
22 a respective object) with respect to this or any other object in this or any other data
23 polyarchy 122. As already noted, these objects can be represented in one or more
24 physically distributed data stores that are in turn logically centralized by one or
25

1 more directory services as well as by one or more data polyarchies. The attribute
2 data field 214 defines any data attributes or data values of the virtual object 210.
3 Each attribute corresponds to the attributes that an actual instance of the virtual
4 object may include. Such attributes include, for example, one or more predicate
5 data fields 216, multiple domain property data fields 218, and zero or more sub-
6 object entity references 220.

7 Each predicate data field 216 indicates a respective operation to access or
8 present a particular object with respect to one or more hierarchies of other objects
9 (each object being represented by virtual objects 210 in the polyarchical data set
10 122). Such operations indicate one or more diverse types of searches (e.g., a linear
11 search and a recursive search), data transformations (e.g., from one hierarchical
12 relationship to another different hierarchical relationship), and so on. (See, block
13 1318 of Fig. 13).

14 If an object is a simple object, meaning that it does not reference to a sub-
15 object entity 220, a predicate 216 operation (e.g., a search, modification, data
16 transformation—from one structure such as from a virtual object 210 to an object
17 within a hierarchy of other objects) will correspond to the respective object of
18 interest. However, if the object is a complex object, meaning that it has a data
19 relationship to one or more sub-objects, then the predicate operation will
20 correspond to a combination of the object and/or the one or more sub-objects.

21 The domain property data field 218 includes a physical domain property
22 and a logical domain property. The physical domain property indicates one or
23 more sets of values used to index a data object. The physical domain property is
24 selected from a group of physical domain properties including a data type, a data
25 precision indication, a scale indication, and a nullability indication). The logical

domain property aspect of the domain property 218 facilitates searching and navigation of the data polyarchy 122 by allowing object data values to be assigned to particular domains. Specifically, the logical domain indicates a strategy to access and/or present the corresponding object with respect to the other objects in the data polyarchy. For instance, the logical domain property includes a unique domain property, a locating domain property, and a classifying domain property. The particular one logical domain property that the polyarchical data relationship management module 116 assigns to an attribute of an object is based on the attribute's relative distribution of its value in the data polyarchy with respect to other values of the same attribute of other objects in the data polyarchy.

We now describe: (a) the relative distribution of the values assumed by an attribute within the data polyarchy 122; and (b) how data distribution determines which objects represent respective dimensions (hierarchies), up-nodes, and down-nodes.

Relative Attribute Value Distribution

The set of values that an attribute has is part of that attribute's logical domain. Any information that is collected about the actual distribution of the values (in terms of the number of potential objects that contain each potential value) in a data polyarchy 111 is also a property of the attribute's logical domain. To determine the relative distribution of attribute values, one or more thresholds (e.g., a low threshold and a high threshold) are defined to determine the attribute's relative distribution in a data polyarchy 122 with respect to other attributes of other objects in the polyarchy. The thresholds are based on the assumption that the data may have a certain percentage of error within it (e.g., one (1) percent

error). (Other statistical analysis techniques can be used in combination with or in place of the thresholds to determine object attribute distributions).

For instance, as objects are loaded into the data polyarchy 122 (e.g., from one or more directory and/or database servers 106), the data polyarchy management module 120 examines each object's respective attributes values based on the thresholds to determine: (a) which attributes are substantially unique with respect to their distributions in objects in the data set; (b) which attributes are distributed across a substantially large set of objects; and (c) which attributes are distributed across a substantially small set of objects in the data set. These determinations are made based on assuming that the data has that certain percentage of error.

With this assumption of some data error in mind, consider that a substantially unique attribute is not necessarily the only attribute of its kind in the data polyarchy 122. Rather, an attribute may be absolutely unique, or the attribute may belong to a relatively sparse distribution of similar attributes in the data set. Attributes that are determined to be substantially unique with respect to their distributions across objects in the data set have a unique logical domain property illustrating that they are distinguishing as compared to other attributes.

Attributes that are distinguishing may identify respective unique dimensions in the polyarchical data set 122, which are represented as up-nodes of an interconnected graph that in turn represents a hierarchical dimension. Inside this model, the default polyarchy is flat. Attributes that are not distinguishing are distributed either across a substantially large set of objects in the data set, or alternatively distributed across a substantially small set of objects. Non-distinguishing attributes are not good candidates for attributes that define

1 dimensions. Instead, such distributions indicate that non-distinguishing attributes
2 belong to one or more of the identified dimensions. Accordingly, a non-
3 distinguishing attribute is represented as a down-node in at least one dimension
4 that is identified by the attributes distribution. Up-node polyarchies are also
5 discovered when all the values of a down-node object are located in a substantially
6 unique up-node object.

7 Attributes that are distributed across a substantially large set of objects have
8 a locating domain property (e.g., a surname may be a locating domain property).
9 Attributes with locating domain properties are used to narrow a search for
10 particular ones of the data objects in the data polyarchy 122. Attributes that are
11 distributed across a substantially small set of objects have a classifying domain
12 property. Attributes with the classifying domain property are used to filter out
13 unwanted objects from a search or navigation procedure.

14 Jump Gates

15 A sub-object entity reference 220 such as a GUID not only indicates
16 whether a virtual object 210 (i.e., a respective object) has a relationship to a
17 different object in the data polyarchy 122, but it also references the different object
18 (i.e., via the different object's corresponding virtual object). Specifically, a sub-
19 object reference uniquely identifies the different object of interest as a sub-object
20 of the virtual object data field. The sub-object reference uniquely identifies the
21 different object of interest across one or more data stores.

22 A virtual object 210 that references a sub-object (via a corresponding sub-
23 object entity reference 220) is a "jump gate". A jump gate represents an elastic
24 data relationship between a complex object and related sub-objects within the
25

polyarchical data set 122. Inter-object data relationships in the data polyarchy are modeled with one or more simple objects 210 and/or complex objects 210. If an object has one or more sub-data relationships, such relationships are either represented as referenced sub-objects 220 in the object (or "surface entity"), or as separate objects 210 linked to another object 210 in some dimension.

To illustrate this, consider that an employee and the employee's dependents are people represented as objects in a directory store. The data store administrators may want to maintain fine-grained information about various aspects of each. To represent sub-world information (about the dependents) in the surface entity (the employee), one can use the following representation shown in TABLE 1.

TABLE 1
EXAMPLE OF STORING SUB-WORLD INFORMATION IN A
SINGLE SURFACE ENTITY

```

<person type="employee" GlueID="13399">
  <name> John Doe </name>
  <age> 31 </age>
  <sex> male </sex>
  <dependents>
    <person type= "spouse">
      <name> Alice Doe </name>
      <age> 31 </age>
      <sex> female </sex>
    </person>
    <person type= "child">
      <name> Sigmund Doe </name>
      <age> 8 </age>
      <sex> male </sex>
    </person>
  </dependents>
  <occupation> forester </occupation>
</person>

```

To represent sub-world information about the dependents in totally distinct entities, Alice Doe and Sigmund Doe would be split off into separate entities, having their own Glue IDs (GUIDs 212), as illustrated, for example, in TABLE 2.

TABLE 2**EXAMPLE OF SEPARATE OBJECT/ENTITY REPRESENTATIONS**

```

<person type= "spouse" GlueID= "24889">
  <relatedEmployee> 13399 </relatedEmployee>
  <name> Alice Doe </name>
  <age> 31 </age>
  <sex> female </sex>
</person>
<person type= "child" GlueID="24890">
  <relatedEmployee> 13399 </relatedEmployee>
  <name> Sigmund Doe </name>
  <age> 8 </age>
  <sex> male </sex>
</person>

```

Note that the "person" elements are identical whether they exist as sub elements in John's virtual entity or as root elements in their own independent virtual entities. In this context, John Doe's entity can be reduced as illustrated in TABLE 3.

TABLE 3**EXAMPLE OF A SINGLE ENTITY REPRESENTATION**

```

<person type="employee" GlueID="13399">
  <name> John Doe </name>
  <age> 31 </age>
  <sex> male </sex>
  <occupation> forester </occupation>
</person>

```

The entity illustrated in TABLE 2 is related to John's dependents along the "dependents" dimension, where "relatedEmployee" is joined to Glue ID to "pass through the jump gate".

Between these two extremes, we can imagine representing John's node internally as illustrated in TABLE 4.

TABLE 4
EXAMPLE OF AN ENTITY REFERENCING ONE OR MORE
OTHER ENTITIES

```

<person type="employee" GlueID="13399">
  <name> John Doe </name>
  <age> 31 </age>
  <sex> male </sex>
  <dependents>
    <person GlueID= "24889"/>
    <person GlueID= "24890"/>
  </dependents>
  <occupation> forester </occupation>
</person>

```

The entity of TABLE 4 could be returned to a client as is allowing the client to add to this information by expanding the related Glue IDs. Or a server such as a data polyarchy server 102 of Fig. 1 could itself de-reference the Glue IDs, returning the following amalgam (shown below in TABLE 5), and demonstrating the elasticity of the solution to the jump gate problem evident in traditional directory implementations.

TABLE 5

EXAMPLE OF DE-REFERENCED IDENTITY INFORMATION

```

<person type="employee" GlueID="13399">
  <name> John Doe </name>
  <age> 31 </age>
  <sex> male </sex>
  <dependents>
    <person type= "spouse" GlueID= "24889">
      <name> Alice Doe </name>
      <age> 31 </age>
      <sex> female </sex>
    </person>
    <person type= "child" GlueID= "24890">
      <name> Sigmund Doe </name>
      <age> 8 </age>
      <sex> male </sex>
    </person>
  </dependents>
  <occupation> archeologist </occupation>
</person>

```

In other words, a virtual object 210 can be modeled as either: (a) a simple object (often referred to as a "simple element") such as a character string, an integer, and so on, that does not reference any other element; or (b) a complex object (often referred to as a "complex element") that references one or more other simple elements or complex elements. In this manner, the polyarchical data set 122 provides for elastic inter-object data relationships that can be defined at any time with any one of a number of different relational representations.

Thus, in sharp contrast to traditional rigid directory implementations that have an intractable schema problem, wherein semantics of shape and naming in a directory must be agreed on in advance to solve the simplest design problem, no fundamental design decision is required when encountering an inter-object data relationship that is modeled as a jump gate. The shape and naming of the directory tree based on the polyarchical data set 122 is not affected by representing

various and elastic inter-object relationships even after a polyarchical data set has been designed. Moreover, an update/modification to a complex object may also result in corresponding updates to one or more related sub-objects that in turn may be represented in one or more different dimensions as compared to a particular dimension that represents the complex object.

Optimizing the Data Polyarchy Schema for De-referenced Operations

Two or more objects can be related to a third object for de-referenced dimensional group, or many-to-many object searching and navigation operations. For example, membership in a group is represented by a membership entity containing information about the relationship between a member and a group. A membership entity includes a *memberOf* data field to identify a group, and a *memberIs* data field to identify a group member. In this implementation such unique identification is accomplished by using respective GUIDs 212.

To determine if an entity is a member of a group, we search for a relationship entity where *memberIs* is the GUID of the entity, and *memberOf* is the GUID of the group. A membership dimension is defined as shown in TABLE 6.

TABLE 6

EXAMPLE OF A MEMBERSHIP DIMENSION IN SCHEMA

```
<dimension dereferenceElement="memberIs">
  <name> membership </name>
  <displayName lang="en"> Membership </displayName>
  <upnodeReferenceElement> memberOf </upnodeReferenceElement>
  <upnodeNamingElement> GlueID </upnodeNamingElement>
</dimension>
```

In this example, the group's GUID (represented in TABLE 6 as "GlueID") identifies the group as an upnode because the GUID is substantially unique, and the children are identified as membership entities with a *memberOf* element set to

1 the group's GUID. A conventional "down" navigation through the data set
2 enumerates the membership entities—which may provide useful information about
3 the nature of each individual membership (e.g. when a particular membership
4 expires).

5 It is also possible to perform an "indirect" enumeration using the *memberIs*
6 association to get information about the actual group members. To do so, issue a
7 "down" enumeration on the group in the membership dimension with de-
8 referencing set to *memberIs*. In this case, the membership entity's *memberIs*
9 element is used to de-reference the actual entity belonging to the group. Thus, it is
10 simple to construct an inverse dimension that list all groups belonged to by an
11 entity. In this case, one may also either list the membership entities, or de-
12 reference them to get information about the groups themselves.

13 Accordingly, no special schema design is required to represent a group's
14 inverse polyarchies or other many-to-many inter-object relationships in the
15 elements-of-interest schema 124.

16 An Exemplary Data Polyarchy Schema

17 Fig. 3 shows further aspects of an exemplary data polyarchy schema 124 of
18 Fig. 1 to indicate how a client can manipulate the data polyarchy 122 in a
19 meaningful manner. The data polyarchy schema is also referred to as an
20 "elements-of-interest" schema. An element is an object attribute or data value.
21 The elements-of-interest schema 124 includes a plurality of data fields 310 to limit
22 a client 110 query on the data polyarchy. Such a query is communicated to the
23 data polyarchy server 102 of Fig. 1. More specifically, such a query is
24
25

1 a dimension or hierarchy (e.g., a GUID, a location, an employee number, a cost
2 center, and so on). The locating index type 310-3 or selecting index type is a good
3 candidate for locating objects within the data polyarchy and is represented, for
4 example, by the following attributes: a surname, a building name, a title, a room
5 number, and/or the like. An attribute having a classifying index type such as an
6 indication of gender (e.g., male or female) is a good candidate to filter objects in a
7 search of objects in the data set because classifying objects are relatively small in
8 number in the data polyarchy as compared to the relative distribution of objects
9 with attributes that correspond to other index types.

10 The elements-of-interest schema 124 is highly customizable. For instance,
11 a network administrator can assign natural language names such as names in
12 English, French, Chinese, etc., to the elements, or objects in the elements-of-
13 interest data set 124. Moreover, the administrator can designate sub-objects for
14 storage as linked but discreet entities, as described in greater detail with respect to
15 jump gates and TABLES 1 through 4. In this manner, objects in the polyarchical
16 data set 122 of Figs. 1 and 2 that would not otherwise be immediately subordinate
17 to a root object become eligible for promotion in the schema. This mechanism is
18 used in conjunction with multiple dimensions (polyarchy) to produce elastic jump
19 gates.

20 TABLE 7 shows an exemplary elements-of-interest schema 124 in an XML
21 data format. Other data format representations besides XML representations (e.g.,
22 an extended version of XML, which has at least a subset or more of the features of
23 XML) of elements 310 are contemplated. In this schema representation, boxed
24 text (i.e., text boxed-in or surrounded with lines) and text preceded by a semi-
25

colon “;” represent corresponding comments. Generally comments of more than a single line are placed in a box.

TABLE 7

EXAMPLE ELEMENTS OF INTEREST SCHEMA

`<WellKnownEntities GlueID="d7a5f1a9-6ba9-48a2-a464-660d82c24b5c">`

; The “WellKnownEntities GlueID” tag is a unique schema ID.

`<ElementsOfInterest>` ; the beginning of the schema

`<element name="objectType">` ; name of the attribute

`<displayName lang="en" value="Object Type"/>`

The “displayname lang” tag indicates a language (“lang”) such as English (“en”) and the corresponding value of the attribute in that language (e.g., cn’s English display name is “Name”). As can be appreciated, the schema can be configured by a network administrator to indicate a number of different display names even for a single attribute (e.g., an English display name, French display name, Chinese display name, and so on).

`</element>`

`<element name="GlueID" indexType="Distinguishing">`

;Note that the “GlueID” element is identified as a “distinguishing” attribute. Other indexTypes include locating, or classifying index types. Each index type is based on the attributes relative distribution in the data polyarchy

`<displayName lang="en" value="Glue ID"/>`

`</element>`

`<element name="cn">`

`<displayName lang="en" value="Name"/>`

`</element>`

`<element name="telephoneNumber">`

`<displayName lang="en" value="Phone Number"/>`

`</element>`

`<element name="roomNumber">`

`<displayName lang="en" value="Room Number"/>`

`</element>`

`<element name="uid">`

`<displayName lang="en" value="E-mail Alias"/>`

`</element>`

`<element name="description">`

`<displayName lang="en" value="Description"/>`

`</element>`

`<element name="sn" indexType="selecting"`

```

1      <displayName lang="en" value="Surname"/>
2  </element>
3      <element          name="givenName"          indexType="locating"
4  startingSize="20000">
5      <displayName lang="en" value="Given Name"/>
6  </element>
7      <element          name="mail"          indexType="distinguishing"
8  indexStartingSize="20000" indexGrowBy="20000">
9      <displayName lang="en" value="E-mail Address"/>
10 </element>
11      <element name="buildingName" indexType="classifying">
12      <displayName lang="en" value="Building Name"/>
13 </element>
14      <element name="title" indexType="classifying">
15      <displayName lang="en" value="Title"/>
16 </element>
17      <element name="location" indexType="distinguishing">
18      <displayName lang="en" value="Location"/>
19 </element>
20      <element name="locationUpnode"/>
21      <element name="uniqueIdentifier" indexType="distinguishing">
22      <displayName lang="en" value="Employee Number"/>
23 </element>
24      <element name="manager">
25      <displayName lang="en" value="Manager"/>
26 </element>
27      <element name="costCenter" indexType="distinguishing">
28      <displayName lang="en" value="Cost Center"/>
29 </element>
30      <element name="costCenterUpnode"/>
31 </ElementsOfInterest>
32 <Dimensions>

```

The "Dimensions" tag is a portion of the schema that identifies those objects in the data polyarchy 122; that represent a root node in a hierarchy of data relationships.

```

<dimension>          ; indicates a dimension
      <name>costCenter</name>          ; name of the dimension

```

```

<upnodeReferenceElement>costCenterUpnode</upnodeReferenceElement>
nt>

```

The "upnodeReferenceElement" tag represents —the element that contains the value that is present in the parent dimension naming element.

`<dimensionNamingElement>costCenter</dimensionNamingElement>`

The "dimensionNamingElement" tag names the objects in that dimension.

`<view>`

The "view" tag can indicate that objects above or below the dimension can be shown (e.g., siblings).

`<displayName lang="en">Business Units</displayName>`

`<SearchType>nodeQuery</SearchType>`

The "SearchType" tag indicates how the client should generate the query. If absent, the client will generate a query that returns a simple list. If "nodeQuery", the generated query will request a hierarchy result in the specified dimension. If "nodeConstraintQuery", the generated query will request a hierarchy constrained to only entities below the specified entity in the specified dimension. If "nodeExclusiveQuery", the generated query will request a hierarchy constrained to be below the first specified entity and not below the second and succeeding specified entities in the specified dimension.

`<up>*</up>`

The "up" tag the number of hierarchy levels (i.e., ancestors in the hierarchy) that are to be displayed in the dimension. In this case, the wild card "*" indicates that all levels in this dimension can be viewed.

`<ElementsList>`

These attributes indicate those that will be displayed at this level. This list is customizable.

`<element>cn</element>`

`<element>uid</element>`

`<element>telephoneNumber</element>`

`<element>title</element>`

`<element>buildingName</element>`

`<element>roomNumber</element>`

`<element>description</element>`

`<element>companyCode</element>`

`<element>costCenter</element>`


```

1      </ElementsList>                                </view>
2      </dimension>
3      <dimension>
4          <name>Management</name>
5
6      <upnodeReferenceElement>manager</upnodeReferenceElement>
7
8      <dimensionNamingElement>uniqueIdentifier</dimensionNamingElement>
9
10     nt>
11
12     <view>
13         <displayNamelang="en">Management</displayName>
14         <displayName lang="fr">Gestion</displayName>
15         <SearchType>nodeQuery</SearchType>
16         <up>*</up>
17         <ElementsList>
18             <element>cn</element>
19             <element>uid</element>
20             <element>telephoneNumber</element>
21             <element>title</element>
22             <element>buildingName</element>
23             <element>roomNumber</element>
24         </ElementsList>
25         <selected>true</selected>

```

The "selected" tag indicates to the client that this view is the default (selected) view in the client interface.

```

16     </view>
17     <view>
18         <displayName                                lang="en">Direct
19         Reports</displayName>
20         <SearchType>nodeQuery</SearchType>
21         <up>0</up>

```

The "<up>0</up>" tag is a dimensional direction indicator that indicates to the client that no hierarchy ancestors of the specified entity should be returned in the specified dimension.

```

22         <down>1</down>

```

The "<down>1</down>" tag is a dimensional direction indicator that indicates to the client that only the immediate descendants (one level of the hierarchy) of the specified entity should be returned in the specified dimension.

```

23         <ElementsList>
24             <element>cn</element>

```

The “SearchElement” tag indicates to the client which element of interest should be queried on the selected entity. For example, if Jane Doe is selected and a “title” searchElement is specified, then the client will determine what Jane Doe’s title is and do a search of all people with that title.

```

1      <ElementsList>
2          <element>cn</element>
3          <element>uid</element>
4          <element>telephoneNumber</element>
5          <element>title</element>
6          <element>buildingName</element>
7          <element>roomNumber</element>
8      </ElementsList>
9  </view>
10 <view>
11     <displayName lang="en">Same Title (list)</displayName>
12     <SearchType>nodeSearch</SearchType>
13     <SearchElement>title</SearchElement>
14     <ElementsList>
15         <element>cn</element>
16         <element>uid</element>
17         <element>telephoneNumber</element>
18         <element>title</element>
19         <element>buildingName</element>
20         <element>roomNumber</element>
21     </ElementsList>
22 </view>
23 </dimension>
24 <dimension>
25     <name>officeLocation</name>
26
27 <upnodeReferenceElement>locationUpnode</upnodeReferenceElement>
28
29 <dimensionNamingElement>location</dimensionNamingElement>
30 <view>
31     <displayName lang="en">Location of
32 Office</displayName>
33     <SearchType>nodeQuery</SearchType>
34     <up>*</up>
35     <ElementsList>
36         <element>cn</element>
37         <element>uid</element>
38         <element>telephoneNumber</element>
39         <element>title</element>
40         <element>buildingName</element>
41         <element>roomNumber</element>
42         <element>description</element>

```

The “Inputs” tag indicates to the Polyarchy data manager where the source material is located and what element should be used as the anchor for that material.

At block 412, responsive to receiving the data (block 410), the data polyarchy management module 120 generates or updates the data polyarchy 122 to reflect any inter-object relationships (e.g., mono-directional and/or bi-directional

relationships) between the received data and the data (if any) already in the polyarchy 122. As already discussed, these inter-object relationships are determined based on the attributes of the received data with respect to the attributes of the other objects in the polyarchy. Specifically, to generate, configure, or update the data polyarchy, the management module analyzes the relative distributions of the attributes of the objects in the data polyarchy to determine which of zero, one, or more dimensions within which each object participates in inter-object relationships with other objects in the polyarchy.

These operations 412 are automatic or dynamic responsive to receipt of the data (block 410) and do not require any intervention of any human operators such as network administrators. Because inter-object relationships in the data polyarchy 122 are determined and expressed based on the values of attributes of the objects in the polyarchy, these inter-object data relationships can be elastic—meaning that they can change over time. As values of attributes change, the inter-object relationships based in the new values are dynamically or automatically represented in the polyarchy by the management module 120 upon receipt. These operations 412 are performed independent of a-priori knowledge of data relationships between respective ones of the data objects in the data polyarchy. Additionally, because inter-object relationships in the data polyarchy are determined and expressed based on the values of attributes of the objects in the polyarchy, these relationships are determined and expressed completely independent of a distinguished name of an object.

At block 414, the data polyarchy management module 120 of Fig. 1 generates, configures, or updates the elements-of-interest schema 124 (e.g., see Figs. 1 and 3) to indicate how the data polyarchy 122 can be manipulated,

presented, and navigated in a meaningful manner. Specifically, as discussed above in reference to Figs. 2 and 3, and Table 7, the schema indicates the elements, or attributes in the data polyarchy along with any corresponding distinguishing, locating, or classifying characteristics of each attribute. The schema also indicates the dimensions in the polyarchy along with each attribute or element of interest contained by objects in the dimension.

An exemplary set of polyarchical query language (PQL) commands (based on the elements-of-interest schema 124) used by a browser to search, navigate, or display portions of the polyarchical data set 122 are described in greater detail below in reference to Figs. 6 through 12. An exemplary procedure to use the elements-of-interest schema 124 to formulate PQL requests and responses is described in greater detail below in reference to Fig. 13.

Exemplary Polyarchical Query Language Request

Fig. 5 shows an exemplary polyarchical query language (PQL) query used by a client 110 to request a data polyarchy server 102 to return information (a PQL response) corresponding to information in the data polyarchy. Responsive to receiving such a query, the data polyarchy management module 120 identifies and retrieves a set of information corresponding to objects in the polyarchy

Queries 500 and corresponding server 102 responses are implemented using a text markup language such as XML. In this configuration, the queries and server responses are packaged in a Simple Object Access Protocol (SOAP) and posted over the network 104 of Fig. 1 using the Hypertext Transfer Protocol (HTTP). SOAP and HTTP are communication protocols that are well known to those skilled in the art of network communication protocols.

1 The message 500 includes a schema ID 502 and one or more object
2 transformation parameters 510 (hereinafter, a parameter is also referred to as a
3 data field) for specifying one or more attributes 214 of Fig. 2. The schema ID is
4 used to identify a particular elements-of-interest schema 124. It can be
5 appreciated that this data field is optional if there is a default schema or only one
6 schema. The attributes 510 correspond to the virtual objects 210 of the data
7 polyarchy 122. (The attribute(s) include distinguishing attributes, locating
8 attributes, or classifying attributes, each of which is discussed in greater detail
9 above with respect to logical domain properties of Fig. 2).

10 A parameter 510, or data field is classified according to its type, which is
11 selected from types that include a specific element of interest type 510-1; an
12 elements-of-interest modifier to limit a response 510-2; a Boolean modifier 510-3;
13 a dimension indicator 510-4; and/or a dimension information modifier 510-5. The
14 number and types of data fields that are represented in the message 500 are based
15 on the message's design, or purpose.

16 Fig. 6 shows a user interface (UI) 600 displaying an exemplary PQL query
17 500 message and a corresponding exemplary PQL response 620. Specifically, the
18 PQL query includes a modifier parameter 510-1 based on a data polyarchy schema
19 124 to specify a particular attribute 510 with which to perform a search of the data
20 polyarchy 122. The UI includes a first area 610 to type in a PQL message 500, a
21 second area 612 to show the PQL message packaged in a SOAP envelope 618 and
22 posted over HTTP, and a third area 614 to show the data polyarchy management
23 module 120 PQL response 620. Although the PQL response is shown as being
24 returned in a SOAP envelope, the response can be returned in a variety of other
25 data packaging formats.

- 1
- 2
- 3
- 4
- 5
- 6
- 7
- 8
- 9

10
11
12
13

14
15
16
17
18
19
20

21

22

23

24

25

1 with respect to polyarchies of data relationships. A Boolean modifier is used to
2 perform a filtering operation (“and”), a union operation (“or”), or an exclusion
3 operation (“not”) on one or more hierarchies of data relationships based on
4 variable. The variable includes an object represented by the data polyarchy
5 schema 124 of Figs. 1 and 3, and Table 7, a hierarchy of objects represented by
6 the schema, or polyarchies of objects represented by the schema, and so on.

7 For example, the “and” Boolean modifier 510-3 is used to filter the results
8 of two data store searches based on specific elements-of-interest data fields 510-1.
9 A first specific elements-of-interest data field specifies a surname (“sn”) attribute
10 with a value of “Smith”. A second specific elements-of-interest data field
11 specifies a “title” attribute with a value of “vice president”. Thus, the Boolean
12 modifier is used to narrow, or filter the results based on the respective search
13 results. The result is a single object in the PQL response 620 that corresponds to
14 vice president John Smith. If there were more than one set of entity information
15 stored in a directory that matched this query 500, then each of the entities would
16 be presented in the result.

17 Fig. 9 shows a user interface 600 that in turn illustrates an exemplary PQL
18 query 500 that includes a dimension information indicator data field 510-4 for
19 specifying a dimension within which to present a response that corresponds to a
20 search operation for an object stored in a data store. In this example, the “under”
21 parameter 510-4 (or “clause”) is combined with a filter 510-3 (“<and>”) to find
22 “architects” under John Smith, which as indicated has a corresponding “unique
23 identifier” of “1234567898”. (See, also John Smith’s unique identifier of Fig. 8).
24 Information corresponding to the architects under John Smith is presented in the
25 PQL response 620 from the server 102. (Note how an elements-of-interest data

1 field 510-2 was used to limit the number of elements presented in the results of the
2 search).

3 Fig. 10 shows user interface 600 for illustrating a PQL query 500 with a
4 dimension information modifier data field 510-5. The dimension information
5 modifier specifies a particular direction and a particular depth to present a data
6 relationship between a complex object in a polyarchical schema and one or more
7 different represented objects. The direction indicates whether the one or more (all
8 objects with the use of a wildcard indication such as "**") different objects are sub-
9 objects of the complex object. The dimension information modifier can also
10 specify SIBLINGS='true' to indicate that all objects with the same parent as the
11 current object should be returned.

12 In this example, the dimension information modifier 510-5 is used to
13 retrieve information 620 corresponding to a first level of subordinates 1010 from a
14 data store. This is a jump gate because John Smith's subordinates 1010 are
15 presented as aspects of John Smith's object definition 620.

16 Fig. 11 is a block diagram of a user interface 600 showing use of a filter
17 parameter in a PQL query 500 with respect to a particular attribute and a
18 subsequent intersection between two polyarchies of data relationships. In this
19 example, two dimensions 510-4 (e.g., a "management" dimension and an "office
20 location" dimension) are intersected and filtered 510-3 based on a "title" attribute
21 of "architect". The search results 620 show the particular objects in the data store
22 that match that query.

23 Fig. 12 shows the user interface 600 for illustrating a PQL 500 that
24 specifies a filter ("and") 510-3 and a union ("or") 510-3 between two
25 polyarchies 510-4 of data relationships. In this example, the filter and the union

are Boolean modifiers. The union attribute is applied to the "management" dimension and the "office location" dimension. The filter specifies a "title" attribute of "architect", which is then applied to the union of the two hierarchies. The search results 620 show the particular objects in the data store that match that query.

Exemplary Procedure to Manage a Polyarchival Data Set

Fig. 13 shows an exemplary procedure 1300 to manage data in a data polyarchy 122. At block 1310, the polyarchival data management module 120 communicates an elements-of-interest schema 124 to a client 110. The elements-of-interest schema 124 indicates to the client how objects in the data polyarchy can be accessed, manipulated, and presented by the client in a meaningful manner.

At block 1312, the polyarchival data management module 120 receives a PQL message 500 that is based on the communicated data polyarchy schema (block 1310). The request not only identifies one or more attributes of interest but also identifies the data relationships of interest. The request corresponds to a data object of the data objects in the polyarchival data set 122 of Fig. 1.

The received PQL message 500 may correspond to one or more operations including any combination of: (a) an operation to find a default search object of the data objects; (b) an operation to locate an object of the data objects that corresponds to a particular name; (c) an operation to obtain a default hierarchy of data relationships that correspond to a particular object of the data objects; (d) an operation to obtain a particular hierarchy of data relationships that correspond to a particular object of the data objects; (e) an operation to identify at least one subset of a plurality of hierarchies of data relationships that correspond to a particular

1 object of the data objects; (f) an operation to obtain multiple hierarchies of data
2 relationships that correspond to a particular object of the data objects; and so on.

3 At block 1314, the data polyarchy management module 120 determines a
4 physical access strategy (e.g., a simple scan, a recursive scan, and so on) to
5 identify data corresponding to the request from the data polyarchy 122. This
6 determination is based on the request (block 1312), which in turn is based on the
7 schema 124 that was communicated to the client 110 (block 1310). As already
8 noted, the schema provides the client not only with information that corresponds
9 to the possible contents of the data polyarchy, but also with includes information
10 describing the possible polyarchies of data relationships that may pertain to any
11 one object of interest (e.g., see the "<Dimension>" indicators shown in Table 7).

12 For instance, consider that if a client request (i.e., a PQL message 500) is
13 designed to filter out all elements-of-interest that pertain to an object with the
14 exception of an absolutely unique distinguishing attribute (e.g., a GUID and a
15 common name that corresponds to the GUID), a simple scan of the data polyarchy
16 122 is an efficient technique to search for information regarding the distinguishing
17 object of interest.

18 The request 500, however, may also indicate that a number of sub-objects
19 should be presented with respect to a complex object (i.e., a jump gate) and then
20 the results are to be subsequently modified by a union of a dimension of
21 information that corresponds to the complex object that is orthogonal to one or
22 more of the sub-objects. In this case, a recursive scan of the data polyarchy 122 is
23 an efficient technique to search for information regarding the objects and inter-
24 object relationships of interest.

1 In this manner, a PQL request message 500 identifying attributes and data
2 relationships of interest also provides an optimized physical access strategy to
3 search the data polyarchy 122 for such attributes and data relationships.

4 At block 1316, the data polyarchy management module 120 accesses the
5 data from the polyarchy based on the determined physical access strategy (block
6 1314). The accessed data may take a number of different forms. For instance, the
7 accessed data may be independent of any inter-object relationship between the
8 data object and any other object in the polyarchy. Additionally, the accessed
9 object(s) may participate in one or more hierarchies of inter-object relationships
10 with one or more different data objects in the polyarchy. In this case, the accessed
11 object(s) and the one or more different objects comprise a similar attribute. As
12 discussed above, these inter-object relationships may be orthogonal with respect to
13 one another in one or more dimensions.

14 At block 1318, the polyarchical data management module 120 transforms
15 the accessed data for issuing to the client 110. Specifically, accessed data is
16 transformed based on the requirements of the specific PQL message 500 that was
17 used to request the data (block 1312). For instance, if the message indicates an
18 object with respect to a particular dimension, the implicit and explicit inter-object
19 relationships of the accessed data are assembled into a hierarchy based on the
20 particular dimension.

21 For example, an accessed data object represents a jump gate when the
22 accessed data includes a complex object of the data objects that is related to one or
23 more sub-objects of the data objects. In this case the complex object is
24 transformed or represented as an independent surface entity. Each of the one or
25 more sub-objects is described as a respective separate entity in a manner that is

1 independent of the surface entity. The one or more sub-objects are then
2 transformed or referenced in the surface entity to indicate a relationship between
3 the complex object and the one or more sub-objects. The referencing is
4 independent of any object naming or hierarchical data relationship between the
5 complex object and the one or more sub-objects.

6 In another example, accessed data includes a first object of the data objects
7 in the polyarchy that is related to one or more sub-objects. The first object is
8 transformed or represented as an independent surface entity. Each of the one or
9 more sub-objects is described as respective separate entities in a manner that is
10 independent of the surface entity. Then, a respective link is included in each of
11 the one or more sub-objects to reference the first object. In this manner, as in the
12 previous example, the data is transformed to express the relationship of interest as
13 indicated in the corresponding PQL message 500.

14 At block 1320, data polyarchy management module 120 issues, or
15 communicates the transformed data (block 1318) to the client.

16 Exemplary Computing Environment

17 Fig. 14 illustrates an example of a suitable computing environment 1400 on
18 which an exemplary data polyarchy server 102 of Fig. 1 may be implemented.
19 The exemplary computing environment is only one example of a suitable
20 computing environment and is not intended to suggest any limitation as to the
21 scope of use or functionality of an exemplary data polyarchy server 102, a
22 server 106, or a client 110. Neither should the computing environment 1400 be
23 interpreted as having any dependency or requirement relating to any one or
24
25

1 combination of components illustrated in the exemplary computing
2 environment 1400.

3 The computer 1402 is operational with numerous other general purpose or
4 special purpose computing system environments or configurations. Examples of
5 well known computing systems, environments, and/or configurations that may be
6 suitable for use with an exemplary computer 1402 include, but are not limited to,
7 personal computers, server computers, thin clients, thick clients, hand-held or
8 laptop devices, multiprocessor systems, microprocessor-based systems, set top
9 boxes, programmable consumer electronics, network PCs, minicomputers,
10 mainframe computers, distributed computing environments that include any of the
11 above systems or devices, and the like.

12 An exemplary computer 1402 may be described in the general context of
13 computer-executable instructions, such as program modules, being executed by a
14 computer. Generally, program modules include routines, programs, objects,
15 components, data structures, and so on, that performs particular tasks or
16 implements particular abstract data types. An exemplary computer 1402 may be
17 practiced in distributed computing environments where tasks are performed by
18 remote processing devices that are linked through a communications network. In
19 a distributed computing environment, program modules may be located in both
20 local and remote computer storage media including memory storage devices.

21 As shown in Fig. 14, the computing environment 1400 includes a general-
22 purpose computing device in the form of a computer 1402. The components of
23 computer 1402 may include, by are not limited to, one or more processors or
24 processing units 1412, a system memory 1414, and a bus 1416 that couples
25

1 various system components including the system memory 1414 to the processor
2 1412.

3 Bus 1416 represents one or more of any of several types of bus structures,
4 including a memory bus or memory controller, a peripheral bus, an accelerated
5 graphics port, and a processor or local bus using any of a variety of bus
6 architectures. By way of example, and not limitation, such architectures include
7 Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA)
8 bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA)
9 local bus, and Peripheral Component Interconnects (PCI) bus also known as
10 Mezzanine bus.

11 Server 1402 typically includes a variety of computer readable media. Such
12 media may be any available media that is accessible by computer 1402, and it
13 includes both volatile and non-volatile media, removable and non-removable
14 media.

15 In Fig. 14, the system memory 1414 includes computer readable media in
16 the form of volatile memory, such as random access memory (RAM) 1420, and/or
17 non-volatile memory, such as read only memory (ROM) 1418. A basic
18 input/output system (BIOS) 1422, containing the basic routines that help to
19 transfer information between elements within computer 1402, such as during start-
20 up, is stored in ROM 1418. RAM 1420 typically contains data and/or program
21 modules that are immediately accessible to and/or presently be operated on by
22 processor 1412.

23 Computer 1402 may further include other removable/non-removable,
24 volatile/non-volatile computer storage media. By way of example only, Fig. 14
25 illustrates a hard disk drive 1424 for reading from and writing to a non-removable,

1 non-volatile magnetic media (not shown and typically called a "hard drive"), a
2 magnetic disk drive 1426 for reading from and writing to a removable, non-
3 volatile magnetic disk 1428 (e.g., a "floppy disk"), and an optical disk drive 1430
4 for reading from or writing to a removable, non-volatile optical disk 1432 such as
5 a CD-ROM, DVD-ROM or other optical media. The hard disk drive 1424,
6 magnetic disk drive 1426, and optical disk drive 1430 are each connected to bus
7 1416 by one or more interfaces 1434.

8 The drives and their associated computer-readable media provide
9 nonvolatile storage of computer readable instructions, data structures, program
10 modules, and other data for computer 1402. Although the exemplary environment
11 described herein employs a hard disk, a removable magnetic disk 1428 and a
12 removable optical disk 1432, it should be appreciated by those skilled in the art
13 that other types of computer readable media which can store data that is accessible
14 by a computer, such as magnetic cassettes, flash memory cards, digital video disks,
15 random access memories (RAMs), read only memories (ROM), and the like, may
16 also be used in the exemplary operating environment.

17 A number of program modules 1440 may be stored on the hard disk,
18 magnetic disk 1428, optical disk 1432, ROM 1418, or RAM 1420, including, by
19 way of example, and not limitation, an operating system 1438, one or more
20 application programs 1440, other program modules 1442, and program data 1444.

21 Each of such operating system 1438, one or more application programs
22 1440 (e.g., a polyarchy data management module 120), other program modules
23 1442, and program data 1444 (e.g., the data polyarchy 122 and the elements-of-
24 interest schema 124)—or some combination thereof, may include an
25

1 implementation of an exemplary data polyarchy server 102 of Fig. 1. Specifically,
2 each may include an implementation of a data polyarchy server 102 to:

- 3 (a) dynamically generate, manage, and update a data polyarchy 122 based on
4 attribute values of the objects;
- 5 (b) analyze the data polyarchy based on relative distribution of attributes to
6 generate an elements of interest schema indicating how objects in the data
7 polyarchy can be meaningfully presented and manipulated within various
8 inter-object relationships;
- 9 (c) communicate the elements-of-interest schema 124 to a client 110;
- 10 (d) responsive to receiving a query (e.g., a PQL message 500) based on the
11 schema, determine a physical access strategy to access the requested data
12 from a polyarchical data set 122;
- 13 (e) access and transform the data based on the query request; and,
- 14 (f) issue the transformed data to the client as a response.

15 A user may enter commands and information into computer 1402 through
16 optional input devices such as keyboard 1446 and pointing device 1448 (such as a
17 "mouse"). Other input devices (not shown) may include a microphone, joystick,
18 game pad, satellite dish, serial port, scanner, or the like. These and other input
19 devices are connected to the processing unit 1412 through a user input interface
20 1450 that is coupled to bus 1416, but may be connected by other interface and bus
21 structures, such as a parallel port, game port, or a universal serial bus (USB).

22 An optional monitor 1452 or other type of display device is also connected
23 to bus 1416 via an interface, such as a video adapter 1454. In addition to the
24 monitor, personal computers typically include other peripheral output devices (not
25

1 shown), such as speakers and printers, which may be connected through output
2 peripheral interface 1455.

3 Computer 1402 may operate in a networked environment using logical
4 connections to one or more remote computers, such as a remote server/computer
5 1462 (e.g., data servers 106). Remote computer 1462 may include many or all of
6 the elements and features described herein relative to computer 1402.

7 Logical connections shown in Fig. 14 are a local area network (LAN) 1457
8 and a general wide area network (WAN) 1459. Such networking environments
9 are commonplace in offices, enterprise-wide computer networks, intranets, and the
10 Internet. When used in a LAN networking environment, the computer 1402 is
11 connected to LAN 1457 via network interface or adapter 1466. When used in a
12 WAN networking environment, the computer typically includes a modem 1458 or
13 other means for establishing communications over the WAN 1459. The modem,
14 which may be internal or external, may be connected to the system bus 1416 via
15 the user input interface 1450 or other appropriate mechanism.

16 Depicted in Fig. 14, is a specific implementation of a WAN via the Internet.
17 Computer 1402 typically includes a modem 1458 or other means for establishing
18 communications over the Internet 1460. Modem, which may be internal or
19 external, is connected to bus 1416 via interface 1450.

20 In a networked environment, program modules depicted relative to the
21 personal computer 1402, or portions thereof, may be stored in a remote memory
22 storage device. By way of example, and not limitation, Fig. 14 illustrates remote
23 application programs 1469 as residing on a memory device of remote computer
24 1462. It will be appreciated that the network connections shown and described are
25

1 exemplary and other means of establishing a communications link between the
2 computers may be used.

3 Computer Readable Media

4 An implementation of an exemplary computer 102 may be stored on or
5 transmitted across some form of computer readable media. Computer readable
6 media can be any available media that can be accessed by a computer. By way of
7 example, and not limitation, computer readable media may comprise "computer
8 storage media" and "communications media."

9 "Computer storage media" include volatile and non-volatile, removable and
10 non-removable media implemented in any method or technology for storage of
11 information such as computer readable instructions, data structures, program
12 modules, or other data. Computer storage media includes, but is not limited to,
13 RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM,
14 digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic
15 tape, magnetic disk storage or other magnetic storage devices, or any other
16 medium which can be used to store the desired information and which can be
17 accessed by a computer.

18 "Communication media" typically embodies computer readable
19 instructions, data structures, program modules, or other data in a modulated data
20 signal, such as carrier wave or other transport mechanism. Communication media
21 also includes any information delivery media.

22 The term "modulated data signal" means a signal that has one or more of its
23 characteristics set or changed in such a manner as to encode information in the
24 signal. By way of example, and not limitation, communication media includes
25

wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, RF, infrared, and other wireless media. Combinations of any of the above are also included within the scope of computer readable media.

Conclusion

The described arrangements and procedures replace traditional notions of distinguished names that represent inter-object relationships within a static hierarchy. More specifically, the described arrangements and procedures replace these traditional notions with dynamically generated graphs of inter-object connections in multiple dimensions of data relationships based on attributes of the objects. In this manner, complex real-world objects are represented with respect to the particular objects themselves, with respect to any set of decomposed sub-entities, or sub-objects that are related to the particular objects. These inter-object relationships are managed and navigated using a data polyarchy schema 124 that has been generated to access elements of interest in the data polyarchy 122.

Although the described subject matter to generate and manage polyarchies of data relationships has been described in language specific to structural features and/or methodological operations, it is to be understood that the subject defined in the appended claims is not necessarily limited to the specific features or operations described. Rather, the specific features and steps are disclosed as preferred forms of implementing the claimed present invention.